ERCAP FY 2008 Request: 82282 **PI:** Edward (Wes) Bethel **Last updated:** 24-OCT-07 (Finished)

SciDAC2 Visualization and Analytics Center for Enabling Technologies

You MUST click the "Save" button at the bottom of each section in order to save your work before s electing another tab

ERCAP Request #82282 for FY2008

1. Principal Investigator

Name	Phone	Email
Bethel, Edward (Wes)	510-486-7353	ewbethel@lbl.gov

2. Other Authorized Preparers

3. Senior Investigators

(Utah) Chris Johnson, Chuck Hansen, Steve Parker
(LLNL) Valerio Pascucci, Hank Childs
(ORNL) Sean Ahern
(UCD) Ken Joy

4. Project Title

SciDAC2 Visualization and Analytics Center for Enabling Technologies

5. Project Name

VACET

6. Project Class

SciDAC

7. Sponsoring Site

USA: Lawrence Berkeley National Laboratory

8. Science Category

Computer Sciences

9. DOE Office and Program

ASCR - Computer Sciences

10. Is this project funded by the DOE Office of Science?

[X] **Yes, this project has direct grant support from the DOE Office of Science**

Who is your DOE Program Manager?

Yukiko Sekine

DOE Office of Science Grant Number(s)

1. SciDAC2 Visualization and Analytics Center for Enabling Technologies (VACET). DOE Grant number N/A as of 10/24/07. Total award amount is $2.2M/yr for five years.
2. High Performance Visualization, MICS Base Program. Award is $500K yr through FY09.

According to LBL Finance people, there are no "Grant Numbers" associated with either of these awar ds but both are "on the books."

11.1 Does this project make use of national security information?

[ ] Yes [X] No

11.2 Does this project collect and store Social Security Numbers, Personally Identifiable Heatlh Information (names or other identifiers matched with health data), Driver's License Numbers, or Financial Account Numbers?

[ ] Yes [X] No

11.3 If this project is given a NERSC award, I agree to monitor the usage associated with it to ensure that, to the best of my ability to determine, usage is for the project described here.

[X] Yes [ ] No

11.4 For continuing projects: I have audited the MPP and/or HPSS usage associated with this project, and to the best of my ability to determine, all usage was for the project specified.

[X] Yes [ ] No

12. Computational Resources Requested

Center	Resource	Alloc Type	Repo	Hours Used 2007	Hours Requested 2008
NERSC	MPP (Usage Charging)	DOE Production	m636	14,251	50,000

13. Mass Storage Resources Requested

Center	Resource	Alloc Type	Repo	SRUs Used 2007	SRUs Requested 2008
NERSC	HPSS	DOE Production	m636	0.0	250,000

14. Justification for Resources Requested

For both the VACET and LBNL Base Program efforts, we anticipate performing several different types of activities to meet our project's objectives. The hours/storage we request reflect our best estimate of the resources needed to accomplish work in these areas:

1. Production visual data analysis: consists of making movies and images for stakeholders.
2. Visual data analysis research and development: we have new projects in the areas of: topologica lly based, temporal feature detection and tracking that will be applied to stakeholder projects in climate and combustion; visual data analysis and scientific dat a management (including machine learning, clustering, PCA/ICA) of large, particle-based datasets for stakeholder projects in fusion and accelerator modeling.

During CY07, we were able to get by using scratch storage and thus didn't consume much of our SRU allocation. We expect this situation to change in CY08, hence the SRU request.

During CY07, our startup allocation of 20K hours worked well. During CY08, we have internally proj ected a substantially greater need that is reflected with our request for 50K CPU hours.

More information about the cycles/storage request:
- The 50K hours is an estimate based upon the following sub-projects:
-- 20K hours for 2D/3D topological analysis work.
-- 20K hours aimed at projects with large, particle-based datasets. Work here will consist if I/O, analysis (machine learning, classification, clustering), visual data anslysis.
-- 10K hours to support work for a broad cross section of projects.
- The 250K SRUs reflect an estimate that is the product of about 50TB in data spread over relative ly large files and over a medium amount of I/O to HPSS.

15. PDSF Repos

16. Project Description

16.1 Project Summary: Provide a brief project description. What will this project accomplish? What is the significance of this work?

Launched in 2006 as one of nine centers under the Department of Energy's Scientific Discovery thro ugh Advanced Computing (SciDAC-2), VACET focuses on leveraging scientific visualization and analytics software technology as an enabling technology fo r increasing scientific productivity and insight. Advances in computational technology have resulted in an information big bang, which in turn has created a sig nificant data understanding challenge. This challenge is widely acknowledged to be one of the primary bottlenecks in contemporary science. The vision for our Cent er is to respond directly to that challenge by adapting, extending, creating when necessary and deploying visualization and data understanding technologies for our sc ience stakeholders at DOE's open computing facilities (NERSC/LBNL and NCCS/ORNL). Using an organizational model as a Visualization and Analytics Center for Enabling Technologies (VACET), we are well positioned to be responsive to the needs of a diverse set of scientific stakeholders, including ot her SciDAC projects, in a coordinated fashion using a range of visualization, mathematics, statistics, computer and computational science and data management tec hnologies.

More specifically, VACET will provide visualization and analytics software infrastructure to suppo rt the challenging data understanding needs of SciDAC Science Applications. This infrastructure will be deployed at DOE's open computing facilities, both NERSC/ LBNL and NCCS/ORNL.

16.2 URL for a relevant web page.

http://www.vacet.org/

17. Accomplishments

1. Visualization of Magneto-rotational instability and turbulent angular momentum transport. Code: VisIt. Objective: Generate movie for customer presentation. Methodology: Compute volume rendered images for various variables and transfer functions. Pick app ropriate one. Render movie of all simulation timesteps. CPU time used: approx. 1k hours. Storage requirements: 450GB for many timesteps of raw data, storage o f intermediate results. Networking requirements (offsite): one transfer of 440G to davinci.nersc.gov. Machines at NERSC be used: davinci.nersc.gov (Altix/SMP). P roject web pages: http://vis.lbl.gov/Vignettes/Incite4/index.html (partial set of images)

2. Comparative Visual Analysis. Several different but related projects aimed to study the relation ships between variables in complex datasets produced by simulation. One of these resulted in a publication in IEEE Transactions on Visualization and Computer Graphics (Proceedings of IEEE Visualization 2007), which is located here: http://vis.lbl.gov/Publications/2007/LBNL-63254-Gosink-IEEEVis07.pdf

3. Remote and distributed visualization. We engaged in numerous direct interactions with science s takeholders to "smooth rough edges" in VisIt's remote/distributed client/server execution model. This type of work typically doesn't result in publications, but is very useful for the customers. The new work is included with the most recent VisIt production release.

4. Data management. Several interrelated projects focus on topics like: (1) improving I/O performa nce into VisIt; (2) solving individual stakeholder data format/conversion issues so they can load their data into VisIt; (3) adding new VisIt database plu gins (file loaders). In all of these smallish projects, the motivation is to perform visual data analysis "close to the data" as the data files tend to be too large to move around.

18. Relevant Publications: Enter in the order authors, title, journal

18.1 Refereed Publications: List all refereed publications in the last 12 months based on research using NERSC resources. You may include publications submitted to journals but not publications in preparation.

October 2007 uJournal Paper. L.J. Gosink, J.C. Anderson, and K.I. Joy, Variable Interactions in Query Driven Visualization, IEEE Transactions on Visualization and Computer Graphics (Proceedings of IEEE Visualization 2007), 13(6), October 2007

June 2007 Journal Paper, E.W. Bethel, C.R. Johnson, K. Joy, S. Ahern, V. Pascucci, H. Childs, J. Cohen, M. Duchaineau, B. Hamann, C. Hansen, D. Laney, P. Lindstrom, J. Meredith, G. Ostrouchov, S.G. Parker, C.T. Silva, A. Sanderson, X. Tricoche. SciDAC Visualization and Analytics Center for Enabling Technology, In Journal of Physics, Conference Series, 78, 2007, Boston MA, USA

June 2007 Journal Paper, Hank Childs. Architectural Challenges and Solutions for Petascale Post processing. Journal of Physics, Conference Series SciDAC 2007, 78, 2007, Boston, MA, USA

June 2007 Journal Paper, C. Jones, K.-L. Ma, A. Sanderson, L. Myers. Visual Interrogation of Gy rokinetic Particle Simulations, In Journal of Physics, Conference Series SciDAC 2007, 78, 2007, Boston MA, USA

June 2007 Conference Proceedings, G. Weber, V. Beckner, H. Childs, T. Ligocki, M. Miller, B. van Straalen, E. W. Bethel, Visualization Tools for Adaptive Mesh Refinement Data. Proceedings of the 4th High End Visualization Workshop, June 2007, Tyrol, A ustria.

18.2 Other Publications: List up to 5 other relevant publications in the last 12 months, also base d on research using NERSC resources.

E. Wes Bethel, Chris Johnson, Cecilia Aragon, Prabhat, Oliver R(ubel, Gunther Weber, Valerio Pascu cci, Hank Childs, Peer-Timo Bremer, Brad Whitlock, Sean Ahern, Jeremey Meredith, George Ostrouchov, Ken Joy, Bernd Hamann, Christoph Garth, Martin Cole, C harles Hansen, Steven Parker, Allen Sanderson, Claudio Silva, Xavier Tricoche. DOE s SciDAC Visualization and Analytics Center for Enabling Technologies Strategy for Petascale Visual Data Analysis Success. To appear in CTWatch Quarterly, Volume 3, Number 4, November 2007.

A. Adelmann, A. Gsell, B. Oswald, T. Schietinger, W. Bethel, J.M. Shalf, C. Siegerist, K. Stocking er. "Progress on H5Part: A Portable High Performance Parallel Data Interface for Electromagnetics Simulations", Particle Accelerator Conference, Albuquerque, Ne w Mexico, USA, June 2007, IEEE Computer Society Press.

19.1 Code and Application Descriptions

Code Name	Description	Mathematics	Numerical Techniques	Machines	Planned Processors	Num Procs Reason
VisIt	Our project will run multiple codes. VisIt is a production-quality, parallel-capable visualization application the uses a client-server architecture for displaying results of visualization to a remote client. It has been used on the BG/L system at LLNL to perform visual data analysis of a grid containing 27 billion zones -- it has proven to be sufficiently scalable on all modern platforms. In addition, it has been ported to run on Jaguar, so we reasonably expect it will run on all platforms at NERSC. Custom topological analysis codes. This family of codes, which are experimental, are run using large, time-varying input. Each timestep of data is processed in serial; parallelization occurs by running many such	For the topological analysis work, please refer to the following Siggraph 2007 and IEEE Visualization 2007 papers authored by a VACET team member: V. Pascucci, G. Scorzelli, P.-T. Bremer, and A. Mascarenhas. "Robust On-line Computation of Reeb Graphs: Simplicity and Speed" In Proceeding of SIGGRAPH 2007 (ACM Transation on graphics), pp. 58-1 to 58-9, 2007. Attila Gyulassy, Vijay Natarajan, Bernd Hamann, Valerio Pascucci. "Efficient Computation of Morse-Smale Complexes for Three-dimensional Scalar Functions" In IEEE Transactions on Visualization and		davinci: 70% franklin: 10% jacquard: 10% bassi: 10%	Typically: 32-64 processors. Hero runs: 256-512 processors.	Computational Requirements, Memory Required, Large number of serial jobs

single-timestep analyses in parallel. Limiting factors for these codes is single-node memory size and I/O bandwidth. For that reason, we expect most of this work to be conducted on davinci.	Computer Graphics, Proceedings of IEEE Visualization 2007, October 2007.	

Particle Methods Linear Solvers Eigensolvers Other Method

Languages	Libraries	Performance Limits	Performance Comments	Checkpoint?	Code Plans
c++,MPI,MPI/IO	hdf4,hdf5,netcdf,OpenGL, Xlib, Qt, Posix communication & I/O, Boxlib, Chombo	MPI I/O, Posix I/O, Single Node Memory Performance, Need large memory footprint for those tasks that are serial	VisIt has successfully run on BG/L at LLNL at approximately 1K-way parallel to do interactive visual data analysis of a 27B zone grid.	N	

19.2 Code and Application Performance

20. Storage and Data Movement

To date, we have been able to get by on using local scratch storage. Peak loads during the past ye ar were on the order of 10-20TB on davinci's scratch filesystem.

In the upcoming year, we anticipate:
- Peak load on davinci's scratch filesystem to be as high as 40-50TB: this level reflects the grow th of size/complexity of our stakeholder's datasets.
- Making better use of NERSC's GFS for projects that need to have data accessible on multiple mach ines; this use model has been, and will likely continue to be, the exception rather than the rule for us during the coming year;
- Best-effort data transfer into/out of NERSC will be fine; multi-stream transfer methods (e.g., p ftp/gridFTP) would be helpful;

21. Data I/O

Statements about typical use model, etc.

- We are better characterized as consumers of data rather than producers of data.
- We'll make a copy of stakeholder's data and place into scratch storage for the duration of the p roject (visual data analysis), which can be days to months.
- Why scratch? Typically, there is no quota and we can have access to multiple 10s of TB of space. Also, at least on davinci, the scratch filesystem has excellent I/O bandwidth. Since visual data analysis is a data intensive activity, we look for the best performin g configuration for our work.
- Copies of data used to achieve results are often stored in HPSS.
- Growth areas in the next year or two will include using high performance indices from the SDM Ce nter (FastBit). This software provides a SQL-like query interface. However, it is *not* a RDBMS like mySQL. There are many papers from the SDM center that show perfo rmance data providing convincing evidence that RDMBS are no good for large, scientific data.
- Our parallel I/O strategy (in VisIt) is flexible: can do collective or non-collective I/O.
- For other projects where a serial code is run concurrently on a bunch of different timesteps, ea ch such process will be doing its own I/O.

22. Analytics Applications

Grid Meshing	File Formats	Data Analysis Applications
	POSIX, MPI IO, HDF5, H5Part, FITS, FSML - fusion simulation markup language (used by several Tech-X codes)	Python, R, VisIt, several custom codes

23. Networking

24. Other HPC Support

A "startup-type" of allocation from NCCS/ORNL to provide support to stakeholders who compute/vis/a nalyze there.

25. Additional Information

Having access to DOE's Open Computing facilities is part of our Center's strategy for delivering a nd deploying visual data analysis software infrastructure to the DOE science community. This idea has strong support within ASCR management.

Our request is likely quite a bit different than the usual request, which likely originate from co mputational science projects who aim to do code development and production runs on large MPP resources. Instead, our project complements those other ones -- we pr ovide visual data analysis capabilities for those other teams. What we share with them is the research, development, testing and production run cycle.

As our project is very much data intensive, we benefit more from machines with lots of I/O capacit y and the largest possible single-process memory footprint.

26. Feedback and Project Requirements

We understand an analytics procurement is underway at NERSC at the present time. We hope the new a nalytics platform will be configured/endowed with capabilities that enable it to be useful for data-intensive activities. Davinci has proven useful because of it s large single-process memory footprint and excellent I/O capabilities. Our effort is,however, hampered by the fact there is no graphics hardware on the floor at NERSC.

In the future, we can foresee the need to have a mechanism for delivering analytics results to our stakeholders via a web-based interface. The alternative is we email images/movies, or make images/movies accessible via world-readable directories. An interesting app roach would be for NERSC to provide project-specific web service (virtual hosting would be fine) whereby authenticated access is granted to authorized user s to access analytics results computed on some NERSC platforms and made available through the web interface. The interconnect might be NGF or some other push/pul l type of mechanism.

27. Current Request Status

- [ ] **Not Finished: This request is not yet ready for review.**
- [x] **Finished: This request is finished and ready for review.**